# ERCOFTAC SIG 45 Workshop

## Uncertainty Quantification and Data Assimilation for Fluid Dynamics

Book of Abstracts

November 11–13, 2026
SISSA - International School for Advanced Studies, Trieste, Italy







ERCOFTAC
European Research Community On Flow, Turbulence And Combustion


# Contents

**Wednesday** **1**

Recent advances in single/multi-fidelity sparse-grids surrogate modeling of parametric PDEs: dealing with noisy evaluations and gradient-enrichment (***Lorenzo Tamellini***) . . . . . 1

Spatio-temporal ROM via spectral proper orthogonal decomposition based Gaussian process regression (***Guglielmo Padula***) . . . . . 2

Diffuse domain FEM solver for the steady Navier–Stokes problem (***Davit Gyulamiryan***) . . . . . 2

Optimal experimental design for thermoacoustic experiments in an industrial test rig (***Matthew Juniper***) . . . . . 3

Learning calibrated forecasts via linear models of non-linear dynamical systems (***Benjamin Kent***) . . . . . 4

Bayesian inverse periodic Navier–Stokes problems for 4D flow MRI reconstruction (***Thomas Cheetham***) . . . . . 4

**Thursday** **5**

Adaptive sampling using generative AI for complex PDEs (***David Silvester***) 5

From turbulence to priors: physics-informed stochastic closure (***William Antolin***) . . . . . 6

PDE-free method for the numerical bifurcation and stability analysis of fluid flows via machine learning (***Alessandro Della Pia***) . . . . . 6

Advected multiplicative diffusion model: a fluid-dynamics-based generative framework (***Valentin Resseguier***) . . . . . 7

Data-assimilation-based calibration of turbulence model constants for the prediction of coherent structures in near-stall airfoil flows (***Antoine Jouin***) . . . . . 7

Uncertainty-aware turbulence closures with Bayesian Kolmogorov–Arnold networks (***Ali Eidi***) . . . . . 8

A data assimilation and machine learning approach for the improved prediction of immersed bodies in oscillatory motion (***Miguel M. Valero***) 9

Localized symbolic regression corrections for RANS turbulence models: a masked SpaRTA approach (***Giovanni Stabile***) . . . . . 9
Bayesian data assimilation and uncertainty quantification of RANS turbulence models in flow MRI (***Claire Namuroy***) . . . . . 10
Towards a generic wall model for streamwise heterogeneous roughness (***Martina Formichetti***) . . . . . 11
Stochastic closure for reduced order models through the combination of Langevin generative modeling and physical knowledge of the system (***Alexis Valls***) . . . . . 11
Bayesian inference of arterial mechanical properties using a one-dimensional blood flow model (***Jose Arnal***) . . . . . 12
Towards digital twins for turbulent flows: incorporating unknown boundary conditions, from synthetic data to experimental validation (***Romain Tiphaigne***) . . . . . 13

**Friday** **14**
DA and ML assisted turbulence modeling for the simulation of turbomachines (***Marcello Meldi***) . . . . . 14
Parametric nonlinear reduced modelling of gas-bearing rotors under geometric degradation (***Giacomo Mossinelli***) . . . . . 15
Neural Markov chain Monte Carlo: likelihood density estimation via normalizing flows and variational autoencoders (***Giacomo Bottacini***) . 15
Bi-fidelity sparse-grid interpolation for uncertainty quantification and sensitivity analysis of complex flows (***Maria Vittoria Salvetti***) . . . . . 16
Graph neural network based non-local models: towards application to RANS turbulence modelling (***Davide Oberto***) . . . . . 17
Robust operator learning for data-driven flow modeling under realistic measurement noise (***Boutheina Tair***) . . . . . 17

# Wednesday

## Recent advances in single/multi-fidelity sparse-grids surrogate modeling of parametric PDEs: dealing with noisy evaluations and gradient-enrichment

Wed 11 Nov
14:30

***Lorenzo Tamellini***

CNR-IMATI, Pavia, Italy

In this talk I will present two recent works aimed at enhancing/robustifying the construction of surrogate models for parametric PDEs with single/multi-fidelity collocation methods, namely Sparse Grids Stochastic Collocation and Multi-Index Stochastic Collocation methods (SGSC / MISC, respectively). The first work concerns robustifying the surrogate model construction with respect to the presence of noise in PDE solvers, due to e.g. iterative method tolerances, pre-asymptotic meshes, etc. The intensity of this corruption might be different in different regions of the parameter domain and hard to estimate and control a priori: SGSC and MISC will end up interpolating such noise, resulting in a spoiled approximation of the response surface, compromising any subsequent analysis (optimization, real-time query, uncertainty quantification). We propose a strategy to detect and prevent such spoiling, based on inspecting the spectral content of the response surface approximation: once the decay of the spectral coefficients stagnates due to noisy evaluations, the construction will either stop (in the single-fidelity scenario) or exclude the corrupted fidelities from that iteration on (in a multi-fidelity scenario). The second work discusses how to incorporate evaluations of gradients of the response function in the construction of SGSC surrogate models. This poses several challenges from a theoretical standpoint, such that eventually it is easier to resort to a least-squares approach where both the samples of the parameter domain and the polynomial space are dictated by a sparse grid construction. We will in particular elaborate on the fact that gradient-enhancement turns out to be not always a convenient operation, depending on a) the relative cost between evaluating the response function and its gradient and b) the expected accuracy of these evaluations.

## Spatio-temporal ROM via spectral proper orthogonal decomposition based Gaussian process regression

Wed 11 Nov 14:55

***Guglielmo Padula***

mathLab, Mathematics Area, SISSA, Trieste, Italy

We propose a spatio-temporal reduced-order modelling strategy that combines Spectral Proper Orthogonal Decomposition and Gaussian Process Regression (SPOD-GPR). GPR provides a non-intrusive probabilistic framework to learn nonlinear input-output maps from data and to quantify predictive uncertainty, making it attractive for uncertainty-aware surrogate modelling. However, when applied to classical POD coefficients, the regression may be affected by the purely energy-based nature of POD, which can mix coherent structures associated with different temporal frequencies. This is particularly limiting for travelling waves, quasi-periodic oscillations, and dynamics with multiple time scales. To address this issue, SPOD is used to extract coherent modes that are optimally ranked in both energy and frequency. In the proposed approach, Gaussian Process models are trained on the complex-valued coefficients of the SPOD expansion associated with each frequency component. This frequency-wise learning strategy provides a more structured representation of the temporal dynamics and improves predictability with respect to standard POD-GPR based on global POD coefficients. The proposed framework is assessed on the two-dimensional Kuramoto–Sivashinsky equation in a quasi-periodic, travelling-wave-like regime and on the fluidic pinball in a quasi-periodic regime.

## Diffuse domain FEM solver for the steady Navier–Stokes problem

Wed 11 Nov 15:20

***Davit Gyulamiryan***

Department of Engineering, University of Cambridge, Cambridge, UK

A FEM solver for the steady Navier–Stokes problem is developed to serve as a forward solver in flow-MRI data assimilation. Quantities such as domain boundary, inlet and outlet BCs are inferred during the assimilation. To avoid refitting a mesh to the varying boundary, diffuse domain method is used. Diffuse domain method approximates the weak form of the original PDE, and offers a straightforward computation of volumetric and surface integrals. This allows fixing the computational mesh during the assimilation and letting the boundary evolve within it, and simplifies the implementation of the assimilation process as a whole. Diffuse domain requires extending boundary conditions from the surface to a thin region near it. The work focuses on Dirichlet BVPs and investigates the solver performance in various approximation regimes. Specifically, different boundary condition extensions, and values of the approximation parameter, $\epsilon$, relative to the mesh size $h$. The developed solver is validated on multiple benchmarks: Beltrami flow in 3D, lid-driven cavity flow

in 2D. As well as compared to an existing forward solver on an *in vitro* data of a flow through an aorta at low and high Reynolds numbers. Additionally, well-posedness and convergence analysis is done for the Poisson problem. The well-posedness of the diffuse approximation is shown in the discrete setting. As well as a priori error estimates are derived and verified with numerical experiments. It is found that in a practical scenario, diffuse domain method yields suboptimal convergence rates – 1.0 for $L^2$ and 0.5 for $H^1$ errors for linear (Q1) elements. Thus, we offer a modification to the forward model which may elevate the performance of the diffuse domain method and match it to exact unfitted-boundary methods, such as the cut-cell.

## Optimal experimental design for thermoacoustic experiments in an industrial test rig

Wed 11 Nov
16:15

***Matthew Juniper***

Department of Engineering, University of Cambridge, Cambridge, UK

Predicting how the thermoacoustic instability of an industrial combustor changes across operating conditions is a long-standing challenge because small uncertainties in the flame response lead to large uncertainties in the thermoacoustic response. We address this challenge for a 1 MegaWatt industrial test rig using a combination of Bayesian inference, Gaussian process regression, and information-theoretic experiment design. Starting from a physics-based model whose parameters are inferred at several operating conditions using Bayesian inference, we use Gaussian process regression to learn how the five parameters of a flame model vary with six operating condition parameters. The resulting Gaussian process model predicts the flame response at unseen conditions with quantified uncertainty, and, when coupled into the acoustic network, produces operating maps of the full thermoacoustic response. We then use metrics from information theory to identify the small number of forcing frequencies that are most informative about the flame parameters. For this industrial test rig, three optimally chosen frequencies recover the flame behaviour to the same fidelity as the full dataset of around 20 frequencies forced from both upstream and downstream, reducing the data required by up to 90%.

## Learning calibrated forecasts via linear models of non-linear dynamical systems

Wed 11 Nov 16:40

***Benjamin Kent***

University of Cambridge, Cambridge, UK

We propose a data-driven approximation technique combining probabilistic modelling and forecasting with dynamics linearisation techniques from Koopman- and Dynamic Mode Decomposition-adjacent literature. Specifically, probabilistic observation functions, a linear dynamics operator and reconstruction operators are constructed. These are optimised to minimise a loss based upon a proper scoring rule, namely the Continuous Ranked Probability Score (CRPS), to directly target a well-calibrated probabilistic approximation. We choose to use CRPS-based optimisation as it is a successful and popular approach, particularly in the weather forecasting community, and its sample-based computation gives great flexibility in the model architecture, and does not need a parametric distribution chosen a priori. The probabilistic forecasts are demonstrated on ODE test problems, and vortex shedding and cavity flow fluid dynamics examples. Results indicate good predictive performance and well-calibrated quantification of uncertainty.

## Bayesian inverse periodic Navier–Stokes problems for 4D flow MRI reconstruction

Wed 11 Nov 17:05

***Thomas Cheetham***

Department of Engineering, University of Cambridge, Cambridge, UK

4D flow MRI provides a non-invasive method to image 3D time-resolved velocity fields of a patient's cardiovascular system, yet long signal acquisition times and limited spatio-temporal resolution restrict its routine clinical adoption. One approach to reducing scan duration to clinically viable levels is to use lower-resolution scans in conjunction with a flow reconstruction algorithm to assimilate the velocimetry data into a CFD model. We formulate an inverse periodic Navier–Stokes problem to infer unknown model parameters by assimilating noisy velocimetry data to reconstruct pulsatile flow fields. By using Gaussian probability distributions within a Bayesian framework, we regularise the inverse problem and obtain the most likely model parameters by maximising the posterior probability. The forward and adjoint Navier–Stokes problems are efficiently solved via the time spectral method to search for periodic solutions in the time domain. To handle the complex, arbitrary cardiovascular geometries which are unknown a priori, we employ a dual-consistent finite-element formulation with an implicit boundary method. The efficacy of this approach is demonstrated for a clinically relevant test case by reconstructing pulsatile flow through a carotid artery from noisy, coarse data.

# Thursday

## Adaptive sampling using generative AI for complex PDEs

Thu 12 Nov
9:45

***David Silvester***

University of Manchester, Manchester, UK

The effectiveness of a flow-based deep generative model for approximating a probability density function from a set of samples is assessed in this talk. Three applications where the strategy will be shown to be useful in a PDE context are the solution of steady-state Fokker–Planck equations, the adaptive generation of collocation points for physics informed neural network approximations, and the generation of new sampling points in the (potentially high-dimensional) parameter space when computing bifurcation boundaries in the assessment of hydrodynamic stability. In this latter application, output from a shallow neural network bifurcation classifier provides a probabilistic measure of flow stability, and the Shannon entropy of these predictions is employed as an uncertainty indicator. This coupling between classification and generative modeling establishes a feedback-driven adaptive learning process analogous to error-indicator based refinement in contemporary PDE solution strategies.

## From turbulence to priors: physics-informed stochastic closure

Thu 12 Nov 10:10

***William Antolin***

INRAE, Paris, France

Geophysical fluid dynamics often relies either on simplified stochastic descriptions or on computationally expensive numerical simulations. We propose a physics-informed stochastic framework that bridges these two approaches by combining stochastic transport with turbulence multiscale decomposition. The objective is to construct realistic ensembles of flow realizations while preserving key physical properties of turbulent flows. In particular, we employ Gaussian Multiplicative Chaos as a generative model for wind fields, enabling the representation of intermittency and extreme small-scale fluctuations that are typically missed by standard Gaussian approaches. The resulting framework provides large ensembles of flow realizations at moderate computational cost. These ensembles can be used for uncertainty quantification and model-error characterization in geophysical fluid mechanics. They also provide physically consistent stochastic priors for data assimilation. More broadly, this framework offers an intermediate path between analytical stochastic models and fully resolved simulations for complex geophysical flows.

## PDE-free method for the numerical bifurcation and stability analysis of fluid flows via machine learning

Thu 12 Nov 10:35

***Alessandro Della Pia***

mathLab, Mathematics Area, SISSA, Trieste, Italy

Inspired by the Equation-Free paradigm, we propose an "embed–learn–lift" framework for the construction of minimal-dimensional reduced-order models (ROMs) for high-fidelity Navier–Stokes simulations, including regimes with symmetries and secondary instabilities that are difficult to capture with standard linear or purely data-driven surrogates. The framework consists of four stages. First, manifold-learning techniques, namely Proper Orthogonal Decomposition (POD) and Diffusion Maps (DMs), are used to identify the intrinsic geometry and dimensionality of the latent manifold underlying the high-dimensional flow dynamics across parameter space. Second, latent-space evolution operators are learned directly from data using Gaussian Process Regression (GPR), yielding probabilistic ROMs with uncertainty quantification (UQ). This allows confidence estimation for the learned dynamics and systematic monitoring of model accuracy as new data are incorporated. Such uncertainty-aware learning is essential because bifurcation and stability analyses are highly sensitive to modelling errors. Third, numerical bifurcation tools are employed directly in latent coordinates to compute bifurcation diagrams and stability properties. The method enables efficient continuation of steady and time-periodic branches arising from Andronov–Hopf and Neimark–Sacker bifurcations, including unstable

limit cycles and Floquet analysis, at a fraction of the cost of full Navier–Stokes simulations. Finally, the pre-image problem is solved to reconstruct the bifurcating states in the original physical space, completing the "lift" step. The methodology is demonstrated on three benchmark flows: cylinder wake, sudden-expansion channel flow, and the fluidic pinball. DMs-based ROMs outperform POD-ROMs by providing geometrically consistent parametrizations and accurately capturing intrinsic dimensionality in the presence of secondary instabilities.

## Advected multiplicative diffusion model: a fluid-dynamics-based generative framework

Thu 12 Nov 11:30

***Valentin Resseguier***

INRAE, OPAALE research unit, Rennes, France

Generative models are increasingly used as tools for data augmentation, uncertainty quantification, and data assimilation. However, classical diffusion models face limitations with intermittency and multiscale images of turbulence. We propose a generative framework based on transport noise: the advected multiplicative score-based generative model (A-MSGM). Like standard score-based diffusion models, A-MSGM has a forward process and a reverse process. Rather than corrupting data by adding noise, the forward process advects and mixes images with a divergence-free Gaussian velocity field. As predicted by MSGM theory, mixed images conserve their energies and evolve over time to weak white noises. Then, the MSGM ELBO provides a score matching loss in order to learn how to revert the mixing process, and eventually generate new turbulence images from white noises. The first numerical experiments on turbulent cylinder wake flow from PIV data demonstrate energy conservation and good generation quality. A-MSGM remains competitive with standard diffusion models and improves the generation accuracy and efficiency.

## Data-assimilation-based calibration of turbulence model constants for the prediction of coherent structures in near-stall airfoil flows

Thu 12 Nov 11:55

***Antoine Jouin***

DAAA, ONERA, Institut Polytechnique de Paris, Meudon, France

A data-assimilation framework combined with linear mean-flow analysis is employed to extract coherent flow fluctuations from mean-flow measurements. Time-averaged PIV data of the flow around a NACA66 airfoil at a chord-based Reynolds number of $Re = 517000$, are used to correct the RANS equations closed with the Spalart–Allmaras model. In contrast to previous approaches introducing an additional forcing term in the turbulence model equation, the present strategy relies on calibrating

the production-related constant $c_{b1}$ of the turbulence model. An optimal value of $c_{b1}^{*} = 0.18$ markedly improves agreement with PIV data in separated-flow conditions, without degrading predictions at lower angles of attack. Global stability and URANS simulations are subsequently performed to evaluate the ability of the correction to better capture the flow dynamics. The results are validated by investigating the low-frequency oscillation of the recirculation bubble, as identified in the SPOD of the PIV measurements.

---

Thu 12 Nov
12:20

# Uncertainty-aware turbulence closures with Bayesian Kolmogorov–Arnold networks

***Ali Eidi***

TU Delft, Delft, The Netherlands

Data-driven turbulence closures have measurably improved Reynolds-averaged Navier–Stokes (RANS) predictions, yet they remain hard to trust: a learned closure is usually opaque, and it returns a single point estimate with no calibrated statement of confidence—even when extrapolating far from its training data. These two weaknesses are typically treated in isolation, but a closure that engineers can rely on has to be both interpretable and uncertainty-aware. We develop a Bayesian tensor-basis Kolmogorov–Arnold network for the Reynolds-stress anisotropy. Learnable univariate spline functions on the network edges make the closure directly inspectable and open to symbolic distillation. Variational posteriors over those splines yield a calibrated predictive distribution with separated epistemic and aleatoric parts. The tensor-basis construction makes predictions invariant, symmetric and traceless by design, and regularises them toward physical realisability. On a controlled regression benchmark we establish the model's calibration and interpretability, and we identify which network parameters should carry the posterior. Applied to Reynolds-stress anisotropy across a hierarchy of flows, the predicted uncertainty is calibrated and informative in distribution, while becoming diagnostically overconfident in extrapolation—a useful, honest signal of where the model leaves familiar territory. Finally, we carry the closure to three-dimensional wind-farm wakes and propagate its posterior through the flow solver, turning parameter uncertainty into quantitative confidence bounds on velocity and turbulence-intensity profiles behind the turbines. Against large-eddy-simulation references on unseen array layouts, the propagated ensemble both corrects the baseline and brackets the reference, demonstrating a closure that improves its mean prediction and reports a trustworthy estimate of its own reliability.

## A data assimilation and machine learning approach for the improved prediction of immersed bodies in oscillatory motion

Thu 12 Nov 12:45

***Miguel M. Valero***

Univ. Lille, CNRS, ONERA, Arts et Métiers ParisTech, Centrale Lille, UMR 9014 – LMFL – Laboratoire de Mécanique des Fluides de Lille – Kampé de Fériet, Lille, France

Computational Fluid Dynamics (CFD) applications are affected by uncertainties arising from physical modeling assumptions, numerical discretization errors, and inaccuracies in the initial and boundary conditions. In recent years, Data Assimilation (DA) tools have proven capable of improving CFD predictions, thereby compensating for the previously mentioned uncertainties. In this work, the DA technique known as the Ensemble Kalman Filter (EnKF) is applied to an oscillating immersed body at $Re = 100$. To do so, some model parameters are set free, and the DA technique, by leveraging sparse high-fidelity measurements in space and time, simultaneously enhances the system state and seeks the best parameter combination that ensures compliance with both model and measurement uncertainties. The enhanced flow fields generated by the DA technique are used to train a Random Forest Regression "black-box" model that learns the corrections applied by the EnKF. With minimal intrusive modifications to the CFD solver, the trained model provides enhanced flow predictions, even under conditions that differ from the training data but exhibit similar physics. To account for the periodic nature of the motion, we develop a sliding-window technique in which multiple models are trained, each specialized for a specific phase of the oscillation cycle. This approach opens the way to studying more complex moving-body flow configurations, where neither the flow physics nor the cycle is known a priori.

## Localized symbolic regression corrections for RANS turbulence models: a masked SpaRTA approach

Thu 12 Nov 14:30

***Giovanni Stabile***

Sant'Anna School of Advanced Studies, Pisa, Italy

The predictive accuracy of RANS turbulence models is fundamentally limited by model-form uncertainty, particularly in separated and adverse-pressure-gradient flows. While data assimilation techniques can leverage high-fidelity data to reduce this uncertainty, existing approaches often sacrifice interpretability, generalization, or implementability. This work introduces a novel regularization strategy within the SpaRTA (Sparse Regression of Turbulent Stress Anisotropy) framework that enhances model robustness and generalization. We propose a spatial filtering approach that restricts symbolic regression training to localized flow regions using masking functions based on the Q-criterion and normalized vorticity. This filtering acts as an implicit regularization mechanism, mitigating overfitting while producing

interpretable algebraic corrections for Reynolds stress anisotropy and turbulent kinetic energy production. Continuous logistic masks ensure smooth transitions, maintaining numerical stability in OpenFOAM implementations. The methodology is validated on Periodic Hills ($\alpha = 1.2$ and $\alpha = 0.8$) and converging-diverging channel flows. Results demonstrate that localized corrections consistently reduce prediction error compared to standard $k$-$\omega$, $k$-$\omega$ SST and the original SpaRTA, with particular improvement in wall forces. Crucially, filtered models show superior generalization to unseen configurations, confirming the regularization effect. Models trained on one geometry transfer effectively to similar flows with different parameters, suggesting reduced sensitivity to training data. This work contributes a practical framework for uncertainty reduction in RANS closures, combining data assimilation through symbolic regression with explicit regularization to improve predictive reliability in engineering applications.

---

Thu 12 Nov
14:55

## Bayesian data assimilation and uncertainty quantification of RANS turbulence models in flow MRI

***Claire Namuroy***

Department of Engineering, University of Cambridge, Cambridge, UK

Flow MRI is a non-invasive medical imaging technique that enables 4D visualisations of fluid flows in the body. Although the method holds great promise for the assessment of cardiovascular disease, the acquired images suffer from noise and low spatial resolution, which prevent direct measurement of subvoxel turbulent flow structures. It is possible, however, to estimate the turbulent kinetic energy (TKE) from the signal loss in flow MRI images. In this study, we use adjoint-accelerated Bayesian inference to jointly assimilate the turbulent flow field and the parameters of zero- and one-equation closure models from noisy flow MRI measurements. We solve an inverse RANS problem with boundary conditions inferred from time-averaged velocimetry and TKE data. We show that the assimilated models successfully reconstruct the mean flow in a confined turbulent jet without overfitting. We subsequently propagate the posterior uncertainty in the model parameters to quantify the uncertainty in the reconstructed flow field. We find that the inferred one-equation models have a higher predictive uncertainty than the algebraic models, and that the inclusion of TKE data further reduces the uncertainty in the reconstruction. The proposed framework provides valuable insights for the development of turbulence models and holds strong potential for patient-specific cardiovascular modelling.

## Towards a generic wall model for streamwise heterogeneous roughness

Thu 12 Nov 15:20

***Martina Formichetti***

University of Southampton, Southampton, UK

This study introduces a novel data-driven framework to investigate the effects of surface roughness on turbulent boundary layers (TBLs). Traditionally, determining the specific properties of rough walls requires either highly detailed physical measurements or computationally expensive high-fidelity simulations, both of which can be impractical for complex applications. To address this challenge, we propose an alternative methodology that combines simplified, smooth-wall baseline RANS simulations with sparse amounts of experimental data. Through a data assimilation process, the framework adjusts key flow variables to indirectly infer the missing surface roughness properties. While this approach successfully captures essential flow behaviour and matches experimental observations for homogeneous roughness, real-world surfaces are rarely uniform. Consequently, this study extends the framework to account for streamwise-varying roughness by integrating sparse experimental measurements with observations on equilibrium recovery. The ultimate goal is to develop a generalised RANS wall model capable of providing accurate friction estimates for heterogeneous surfaces, offering a highly efficient and practical tool for complex aerodynamic applications.

## Stochastic closure for reduced order models through the combination of Langevin generative modeling and physical knowledge of the system

Thu 12 Nov 16:15

***Alexis Valls***

Inria, France

Monitoring and control of turbulent fluid systems remain major challenges in many applications, including environmental flows and biomedical systems such as hemodynamics. These applications rely on data assimilation to integrate sparse and noisy observations into predictive models, which requires real-time state estimation and forecasting. In this context, high-fidelity numerical simulations become computationally prohibitive, making reduced-order modeling (ROM) a key ingredient for tractable real-time analysis and control. In this work, we propose a stochastic closure approach for ROMs of chaotic fluid dynamical systems. The approach constrains the ROM to reproduce the statistical behavior of a corresponding high-fidelity simulation while preserving a physics-informed drift structure derived from the underlying POD–Galerkin formulation. The resulting reduced dynamics naturally take the form of a stochastic differential equation, akin to a Langevin process. However, unlike standard generative models (e.g., Langevin dynamics, diffusion models and normalizing flows), our ROM encodes not only the attractor but also the physics-based dynamics on

the attractor. This formulation provides intrinsic uncertainty quantification capabilities and offers a natural foundation for data assimilation strategies. The method is assessed on several dynamical regimes of the Kuramoto–Sivashinsky equation. Numerical experiments show that the proposed approach successfully reproduces both long-term statistical properties and dynamical behavior of the reference system.

---

Thu 12 Nov
16:40

## Bayesian inference of arterial mechanical properties using a one-dimensional blood flow model

***Jose Arnal***

University of Cambridge, Cambridge, UK

Arteriosclerosis—increased arterial stiffness—is a well-established predictor of cardiovascular risk, with loss of vascular compliance linked to the onset of heart failure, stroke, and coronary artery disease. The accurate characterisation of arterial mechanical properties is therefore of significant clinical interest; however, patient-specific, non-invasive estimation of vascular compliance remains an open challenge. To this end, we formulate a Bayesian inverse problem in which haemodynamic data is assimilated into a one-dimensional compliant-vessel blood flow model to infer patient-specific arterial mechanical properties. The forward model is derived from the Navier–Stokes equations under geometric and velocity-profile assumptions, coupled with an elastic constitutive relation between transmural pressure and vessel wall displacement. We discretise the governing partial differential equations in space with a finite-element method and solve for periodic flows in time with a time-spectral approach. The inverse problem assumes Gaussian error statistics and is solved efficiently via adjoint-based gradient evaluation. Additionally, we employ the Laplace approximation to quantify posterior uncertainties in the inferred mechanical properties. The proposed framework is validated on synthetic and experimental data, demonstrating the potential of data assimilation as a tool for patient-specific cardiovascular diagnosis.

# Towards digital twins for turbulent flows: incorporating unknown boundary conditions, from synthetic data to experimental validation

Thu 12 Nov
17:05

***Romain Tiphaigne***

INRAE, Paris, France

Coupling model reduction with limited observations using data assimilation enables digital twins for turbulent flows. To address advection-dominated systems, the reduced-order model (ROM) requires closure. RedLUM achieves this via stochastic closure, modeling unresolved velocity with spatially correlated transport noise. ROM predictions are further improved by assimilating measurements with a particle filter. Recently, we got a step closer to predicting real-world flows. For higher Reynolds numbers, we relied on simulations with a turbulence closure model. We implemented hyperreduction to handle the nonlinearities arising from the turbulence model within the ROM and proposed a new error quantification method for this additional approximation, thereby improving the ROM's accuracy. A key advancement in the present study is the explicit introduction of unknown turbulent boundary conditions. We address these by incorporating complementary boundary modes, which enable the ROM to capture realistic, unmeasured boundary effects. The dynamics of the traditional ROM modes are thus coupled with those of the complementary modes, enhancing the model's ability to represent turbulent flows with uncertain boundaries. We will incorporate results from sparse, unsteady wind-tunnel experiments at a Reynolds number of 3900, thereby advancing beyond our previous work, which relied solely on synthetic data. Looking ahead, we aim to address even higher Reynolds numbers, targeting simulations up to $10^6$.

# Friday

## DA and ML assisted turbulence modeling for the simulation of turbomachines

Fri 13 Nov
9:45

***Marcello Meldi***

ENSAM, Paris, France

The development and assessment of a hybrid framework coupling numerical simulation, data assimilation (DA) and machine learning (ML) to investigate complex internal flows within an axial compressor stage is presented. The methodology first implements a Multi-Fidelity Ensemble Kalman Filter (MFEnKF) that integrates experimental PIV data with a two-fidelity numerical simulation setup, combining steady-state RANS on fixed grid with unsteady hybrid RANS/LES simulations on a moving grid. While this data assimilation step successfully reduces discrepancies between experiments and CFD realizations, its high computational demand limits its direct applicability to industrial geometries. To bypass these computational bottlenecks, machine learning (ML) is used to develop surrogate models trained on the assimilated flow fields. The neural network architecture maps local physical features such as velocity, gradients and viscosity, to predict forcing terms in the Navier–Stokes equations which mimic the data assimilation update. The surrogate model dramatically reduces the inference time while maintaining predictive accuracy comparable to the experimental uncertainty across the studied domain. It is however seen that localized errors in specific regions of the compressor, such as close to the blade tips, can exceed those from the baseline RANS simulations. To resolve this issue, classification of flow regions and combination of different ML models is used to improve the local accuracy of the resulting flow prediction.

## Parametric nonlinear reduced modelling of gas-bearing rotors under geometric degradation

Fri 13 Nov
10:10

***Giacomo Mossinelli***

Institute of Mathematics, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

Rotor–gas-bearing systems couple a nonlinear Reynolds-type lubrication equation for compressible gas-film pressure with rigid-body equations for rotor motion. Their predictive simulation is sensitive to geometry-induced perturbations: dust deposition, groove deterioration or local clearance variations enter the coefficients of the homogenized gas-film model, modify the pressure field and affect the resulting forces and stability properties. A parametric nonlinear reduced-order formulation is considered for pressure-driven rotor dynamics. Pressure snapshots are compressed by POD to separate a few active modes from higher-order components; the latter are represented as a learned graph of the active coordinates, defining a nonlinear trial manifold for the pressure field. ECSW hyper-reduction reduces nonlinear residual and Jacobian evaluations, while targeting the force accuracy and stability of larger linear ROMs at lower online cost. Geometric quantities are treated as parameters of the pROM. Beyond reproducing the dynamics, the reduced coordinates provide a compact state in which macroscopic measurements, such as rotor position, can be used to estimate effective geometric parameters in the homogenized model. Updating these parameters from observations can correct the simulated response and support stability assessment and control strategies aimed at avoiding anomalous operating regimes.

## Neural Markov chain Monte Carlo: likelihood density estimation via normalizing flows and variational autoencoders

Fri 13 Nov
10:35

***Giacomo Bottacini***

MOX, Dipartimento di Matematica, Politecnico di Milano, Milan, Italy

In this talk we present a Bayesian framework that combines Markov chain Monte Carlo (MCMC) sampling, dimensionality reduction, and neural density estimation to efficiently solve inverse problems characterized by intractable or unavailable likelihood functions. The posterior probability distribution over quantities of interest is estimated via differential evolution Metropolis sampling, empowered by learnable mappings. First, a variational autoencoder performs probabilistic feature extraction from observational data. The resulting latent structure inherently quantifies uncertainty, capturing deviations between the actual data-generating process and the training data distribution. At each step of the MCMC random walk, the algorithm jointly samples from the data-informed latent distribution and the space of parameters to be inferred. These samples are fed into a neural likelihood estimator based

on normalizing flows, specifically real-valued non-volume preserving transformations. The scaling and translation functions of the affine coupling layers are modeled by neural networks conditioned on the unknown parameters, allowing the representation of arbitrary observation likelihoods. The proposed methodology is validated on two case studies: (i) the structural health monitoring of a railway bridge for damage detection, localization, and quantification, and (ii) the estimation of the conductivity field in a steady-state Darcy's groundwater flow problem. The results demonstrate the efficiency and accuracy of the inference strategy, while ensuring that model-reality mismatches do not result in overconfident, yet inaccurate, estimates.

## Bi-fidelity sparse-grid interpolation for uncertainty quantification and sensitivity analysis of complex flows

Fri 13 Nov
11:30

***Maria Vittoria Salvetti***

University of Pisa, Pisa, Italy

High-fidelity computational models of complex flows provide accurate predictions, but their significant computational cost often makes them impractical for applications requiring repeated evaluations like, e.g., uncertainty quantification or systematic parametric studies for reliability assessment or optimization problems. To overcome this limitation, surrogate modeling has emerged as a strategy to approximate the mapping between input parameters and output responses in situations where evaluating the "true" relationship via high-fidelity simulations is computationally prohibitive. Among the available approaches, stochastic collocation based on sparse grids has become a widely adopted method. However, they are known to be inefficient for problems requiring local refinement, such as when the target function exhibits localized features or sharp gradients. In this work we present a strategy for improving the local efficiency of Lagrangian sparse grids. We do so by building the sparse grid approximation incrementally and evaluating the function only at collocation points at which a suitable (and crucially, zero-cost) error indicator suggests that incorporating the function evaluation would significantly change the landscape of the approximation. The methodology is applied to the sensitivity analysis to simulation parameters in two challenging flow problems: the numerical prediction of flashback phenomena in hydrogen-fueled perforated burners and the simulation of rogue-wave dynamics in the ocean.

## Graph neural network based non-local models: towards application to RANS turbulence modelling

Fri 13 Nov 11:55

***Davide Oberto***

mathLab, Mathematics Area, SISSA, Trieste, Italy

The closure problem of the Reynolds-averaged Navier–Stokes (RANS) equations remains one of the main challenges in turbulence modelling. In recent years, machine learning (ML) has emerged as a promising alternative to traditional heuristic modelling approaches. While turbulence is intrinsically a non-local phenomenon governed by diffusive and transport processes, the majority of ML-based RANS models are local. Specifically, ML-based models usually predict the Quantity of Interest (QoI) at a specific spatial point using only the input values at the same point, neglecting the surrounding neighbourhood. Only a limited number of studies have incorporated non-local effects, giving very promising results. However, such models need non-trivial input transformations to guarantee frame-independency of the model. In this talk, we explore the use of graph neural networks (GNNs) as ML framework to close RANS equations. GNNs exploit the natural correspondence between computational meshes used in CFD simulations and undirected graphs, enabling automatically the enforcement of frame-independency without any specific input transformation. We analyse GNNs effectiveness and compare them with standard feed-forward neural networks (FNNs) using the classic benchmark periodic hill flow.

## Robust operator learning for data-driven flow modeling under realistic measurement noise

Fri 13 Nov 12:20

***Boutheina Tair***

École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

Data-driven surrogate models offer a promising route toward accelerating computational fluid dynamics (CFD), with potential applications in real-time prediction, digital twins, and data assimilation. However, their practical deployment is often limited by the strong assumptions made about the quality of training and observational data. In particular, flow measurements can exhibit heavy-tailed, spatially or temporally correlated, and otherwise structured noise, for which conventional pointwise regression objectives can be highly sensitive. In this talk, we present a robust operator-learning framework designed to learn mappings between function spaces directly from imperfect observations. Rather than enforcing a pointwise correspondence between predictions and noisy measurements, we formulate training as a distributional weak matching problem: empirical distributions of predicted and observed fields are matched within mini-batches using probability metrics such as total variation and optimal transport. This relaxes the sample-wise correspondence while preserving the computational advantages of neural operators. The approach therefore provides a natural data-driven mechanism for incorporating uncertainty

in measurements directly into the learning procedure, without requiring a separate denoising or filtering stage. We focus in particular on the learning of operators associated with Navier–Stokes flows, and investigate their robustness under increasingly severe measurement corruption. Experiments with heavy-tailed Cauchy noise show that distributional weak matching can substantially outperform conventional pointwise objectives at high noise levels, while recovering qualitatively more accurate flow structures, such as an improved reconstruction of vorticity under strong corruption. Finally, we discuss a perturbation analysis that characterizes robustness in terms of the geometry of the representation, the sensitivity of the distributional discrepancy, and the covariance structure of the measurement perturbations. This work highlights how robust operator learning can provide an efficient bridge between CFD, noisy experimental data, and uncertainty-aware data-driven modeling, with the potential to support reliable surrogate modeling and future integration with data assimilation and flow-control pipelines.